

Harmonic Singular Integrals and Steerable Wavelets in $L_2(\mathbb{R}^d)$ *

John Paul Ward [†] Michael Unser [‡]

February 25, 2013

Abstract

Here we present a method of constructing steerable wavelet frames in $L_2(\mathbb{R}^d)$ that generalizes and unifies previous approaches, including Simoncelli's pyramid and Riesz wavelets. The motivation for steerable wavelets is the need to more accurately account for the orientation of data. Such wavelets can be constructed by decomposing an isotropic mother wavelet into a finite collection of oriented mother wavelets. The key to this construction is that the angular decomposition is an isometry, whereby the new collection of wavelets maintains the frame bounds of the original one. The general method that we propose here is based on partitions of unity involving spherical harmonics. A fundamental aspect of this construction is that Fourier multipliers composed of spherical harmonics correspond to singular integrals in the spatial domain. Such transforms have been studied extensively in the field of harmonic analysis, and we take advantage of this wealth of knowledge to make the proposed construction practically feasible and computationally efficient.

1 Introduction

Building upon [27], our purpose in this paper is to provide a systematic and practical approach to the construction and implementation of steerable wavelets in higher dimensions. The basis of this construction is the theory of

*This research was funded in part by ERC Grant ERC-2010-AdG 267439-FUN-SP.

[†]Biomedical Imaging Group, École polytechnique fédérale de Lausanne (EPFL), Station 17, CH-1015, Lausanne, Switzerland (john.ward@epfl.ch).

[‡]Biomedical Imaging Group, École polytechnique fédérale de Lausanne (EPFL), Station 17, CH-1015, Lausanne, Switzerland (michael.unser@epfl.ch).

singular integral transforms on \mathbb{R}^d with kernels of the form $\Omega(\mathbf{x}/|\mathbf{x}|)/|\mathbf{x}|^d$, where Ω is a smooth function defined on \mathbb{S}^{d-1} . Properties of such transforms were studied by Mikhlin [13], and Calderón and Zygmund [3]. Useful resources for this material are the books of Stein and Weiss, [24, 26], which we shall use as primary references. An attractive feature of singular integral transforms is their correspondence with Fourier multiplier transforms. For instance, if Ω is a spherical harmonic, then the singular integral transform corresponds to a Fourier multiplier that is a multiple of Ω . Furthermore, the spherical harmonics (in particular, the zonal spherical harmonics) satisfy symmetry properties which make them ideal for applications requiring rotations.

The two key ingredients for our construction of steerable wavelets are:

- 1) an isotropic, band-limited mother wavelet ψ that generates a primary wavelet frame of $L_2(\mathbb{R}^d)$;
- 2) a finite collection of functions $\{m_n\}_{n=1}^{n_{\max}}$ which generate a partition of unity on the sphere:

$$\sum_{n=1}^{n_{\max}} |m_n(\boldsymbol{\omega})|^2 = 1,$$

where the m_n are purely polar functions; i.e., $m_n(\boldsymbol{\omega}) = m_n(\boldsymbol{\omega}/|\boldsymbol{\omega}|)$.

The steerable wavelet frame is then generated by the functions $\{\mathcal{F}^{-1}\{m_n\widehat{\psi}\}\}$. The limitation on the functions m_n are minimal; however, we shall focus on zonal spherical harmonics as they make implementation more amenable. Notice that decomposing a signal in this enlarged dictionary of wavelets provides more information about the local orientation of the data; meanwhile, the partition of unity property guarantees that the frame bounds are preserved.

We shall devote the remainder of this section to some basic notation. In Section 2, we shall recall some results about singular integral transforms and their relation to tight frames. In Section 3, we shall cover some details about spherical harmonics and properties of related transforms. Finally, in Section 4, we shall describe the steerable wavelet construction, and in Section 5 we conclude with some specifics about wavelets based on spherical harmonics.

1.1 Notation

The function spaces that we shall consider are the $L_p(\mathbb{R}^d)$ spaces, with norm

$$\|f\|_{L_p(\mathbb{R}^d)} = \left(\int_{\mathbb{R}^d} |f(\mathbf{x})|^p d\mathbf{x} \right)^{1/p},$$

for $1 \leq p < \infty$, and our primary focus shall be with $L_2(\mathbb{R}^d)$. Following the notation of [26], we define the Fourier transform of a function $f \in L_1(\mathbb{R}^d)$ to be

$$\begin{aligned}\widehat{f}(\boldsymbol{\omega}) &= \mathcal{F}\{f\}(\boldsymbol{\omega}) = \int_{\mathbb{R}^d} f(\mathbf{x}) \overline{e^{2\pi i \mathbf{x} \cdot \boldsymbol{\omega}}} d\mathbf{x} \\ &= \int_{\mathbb{R}^d} f(\mathbf{x}) e^{-2\pi i \mathbf{x} \cdot \boldsymbol{\omega}} d\mathbf{x}\end{aligned}$$

and the inverse Fourier transform of f is denoted by $\mathcal{F}^{-1}\{f\}$. For a radial function $f(\mathbf{x}) = f_r(|\mathbf{x}|)$, we can write the Fourier transform as

$$\begin{aligned}\mathcal{F}\{f\}(\boldsymbol{\omega}) &= \int_0^\infty t^{d-1} f_r(t) \int_{\mathbb{S}^{d-1}} e^{-2\pi i t(\mathbf{x}/|\mathbf{x}|) \cdot \boldsymbol{\omega}} d\sigma\left(\frac{\mathbf{x}}{|\mathbf{x}|}\right) dt \\ &= 2\pi |\boldsymbol{\omega}|^{-(d-2)/2} \int_0^\infty f_r(t) J_{(d-2)/2}(2\pi |\boldsymbol{\omega}| t) t^{d/2} dt\end{aligned}$$

where σ is the usual surface measure on \mathbb{S}^{d-1} and $J_{(d-2)/2}$ is the Bessel function of the first kind of order $(d-2)/2$, cf. [25, Section VIII.3]. The area of the sphere is given by

$$\sigma(\mathbb{S}^{d-1}) = \frac{2\pi^{d/2}}{\Gamma(d/2)},$$

where Γ denotes the Gamma function.

For vector valued functions $f : \mathbb{R}^d \rightarrow \mathbb{C}^N$, we denote the n th component by $[f]_n$. The space of such functions, all of whose components are $L_2(\mathbb{R}^d)$ functions, will be denoted by $L_2^N(\mathbb{R}^d)$, and we define the norm

$$\|f\|_{L_2^N(\mathbb{R}^d)} = \left(\sum_{j=1}^N \|[f]_j\|_{L_2(\mathbb{R}^d)}^2 \right)^{1/2}.$$

2 Singular integrals and Fourier multipliers

In this section, we recall some relevant results from the theory of singular integrals and provide a basis for the construction of steerable wavelets. One of the key ingredients in this construction is a collection of functions which generate a partition of unity. In this section we shall show how particular classes of such functions behave as Fourier multipliers.

Definition 2.1. A collection of complex valued functions $\mathcal{M} = \{m_n\}_{n=1}^{n_{\max}}$ will be called admissible if

1. Each m_n is Lebesgue measurable and homogeneous of degree 0; i.e., $m_n(a\boldsymbol{\omega}) = a^0 m_n(\boldsymbol{\omega}) = m_n(\boldsymbol{\omega})$ for all $a > 0$ and $\boldsymbol{\omega} \neq 0$;
2. The squared moduli of the elements of \mathcal{M} form a partition of unity:

$$\sum_{n=1}^{n_{\max}} |m_n(\boldsymbol{\omega})|^2 = 1$$

for every $\boldsymbol{\omega} \in \mathbb{R}^d \setminus \{\mathbf{0}\}$.

The partition of unity property implies that $|m_n(\boldsymbol{\omega})|^2 \leq 1$, so each function is a valid Fourier multiplier on $L_2(\mathbb{R}^d)$; i.e.

$$\left\| \mathcal{F}^{-1}\{m_n \widehat{f}\} \right\|_{L_2(\mathbb{R}^d)} \leq \|f\|_{L_2(\mathbb{R}^d)}.$$

Therefore, we can define a transform mapping $L_2(\mathbb{R}^d)$ to the vector valued space $L_2^{n_{\max}}(\mathbb{R}^d)$ as follows.

Definition 2.2. Given an admissible collection \mathcal{M} , define the transform $T_{\mathcal{M}} : L_2(\mathbb{R}^d) \rightarrow L_2^{n_{\max}}(\mathbb{R}^d)$ by

$$[T_{\mathcal{M}}(f)]_n = \mathcal{F}^{-1}\{m_n \widehat{f}\},$$

and its adjoint $T_{\mathcal{M}}^* : L_2^{n_{\max}}(\mathbb{R}^d) \rightarrow L_2(\mathbb{R}^d)$ by

$$T_{\mathcal{M}}^*(f) = \mathcal{F}^{-1} \left\{ \sum_n \overline{m_n} [\widehat{f}]_n \right\}.$$

Note that, since we apply these transforms to wavelets, the homogeneity condition makes sense, as it means that m_n is invariant to scaling.

Property 2.3. The transform $T_{\mathcal{M}}$ maps a wavelet family into another one in the sense that

$$[T_{\mathcal{M}}(\psi(\cdot/a - \mathbf{b}))]_n(\mathbf{x}) = [T_{\mathcal{M}}(\psi)]_n(\mathbf{x}/a - \mathbf{b})$$

for any $\psi \in L_2(\mathbb{R}^d)$, $a \in \mathbb{R}^+$, and $\mathbf{b} \in \mathbb{R}^d$.

Also, it follows from Plancherel's identity and the partition of unity condition that $T_{\mathcal{M}}$ is in fact an isometry. Hence, we can apply this transform to a tight frame to generate a new frame with the same frame bounds.

Theorem 2.4. Suppose $\{\phi_k : k \in \mathbb{Z}\}$ is a Parseval frame for $L_2(\mathbb{R}^d)$; i.e.,

$$f = \sum_k \langle f, \phi_k \rangle \phi_k, \quad (1)$$

and

$$\sum_k |\langle f, \phi_k \rangle|^2 = \|f\|_2^2 \quad (2)$$

for every $f \in L_2(\mathbb{R}^d)$. Then the multiplier transform $T_{\mathcal{M}}$ associated with an admissible collection \mathcal{M} generates a related Parseval frame:

$$\{\psi_{k,n} = [T_{\mathcal{M}}(\phi_k)]_n : k \in \mathbb{Z}, n = 1, \dots, n_{\max}\},$$

with

$$f = \sum_{n=1}^{n_{\max}} \sum_k \langle f, \psi_{k,n} \rangle \psi_{k,n}.$$

Proof. Our proof follows the same lines as [29, Proposition 1]. Notice that if \mathcal{M} is admissible, then so is $\overline{\mathcal{M}} := \{\overline{m}_n\}_{n=1}^{n_{\max}}$. From the definition of \mathcal{M}

$$f = T_{\overline{\mathcal{M}}}^* T_{\mathcal{M}} f.$$

Additionally, each component of $T_{\overline{\mathcal{M}}} f$ can be expanded in the original frame

$$f = T_{\overline{\mathcal{M}}}^* \mathbf{F}$$

where \mathbf{F} is the function with components

$$\begin{aligned} [\mathbf{F}]_n &= \sum_k \langle [T_{\overline{\mathcal{M}}}(f)]_n, \phi_k \rangle \phi_k \\ &= \sum_k \langle f, [T_{\mathcal{M}}(\phi_k)]_n \rangle \phi_k \end{aligned}$$

The reproduction property now follows by computing the product $T_{\overline{\mathcal{M}}}^* \mathbf{F}$:

$$\begin{aligned} f &= \mathcal{F}^{-1} \left\{ \sum_n m_n \sum_k \langle f, [T_{\mathcal{M}}(\phi_k)]_n \rangle \widehat{\phi}_k \right\} \\ &= \sum_n \sum_k \langle f, \psi_{k,n} \rangle \psi_{k,n}. \end{aligned}$$

To verify that the frame is still tight, write

$$\begin{aligned}\| [T_{\overline{\mathcal{M}}}(f)]_n \|_2^2 &= \sum_k |\langle [T_{\overline{\mathcal{M}}}(f)]_n, \phi_k \rangle|^2 \\ &= \sum_k |\langle f, \psi_{k,n} \rangle|^2,\end{aligned}$$

so that summing over n gives the result. \square

Notice that our Definition 2.1 of admissibility is fairly general and directly exploitable for implementing wavelets in the Fourier domain. However, if a feasible spatial domain representation is required, we must impose certain restrictions. A reasonable condition is to assume that the elements of \mathcal{M} are smooth, since it allows us to relate the multiplier transform to a singular integral transform.

Theorem 2.5. *[24, Theorem III.6] Let m be homogeneous of degree 0 and indefinitely differentiable on \mathbb{S}^{d-1} . Then for $1 < p < \infty$ the Fourier multiplier transform $T : L_p(\mathbb{R}^d) \rightarrow L_p(\mathbb{R}^d)$ given by*

$$T(f) = \mathcal{F}^{-1}\{\widehat{mf}\}$$

can be computed by the singular integral

$$T(f)(\mathbf{x}) = cf(\mathbf{x}) + \lim_{\epsilon \rightarrow 0} \int_{|\mathbf{y}| > \epsilon} \frac{\Omega(\mathbf{y})}{|\mathbf{y}|^d} f(\mathbf{x} - \mathbf{y}) d\mathbf{y},$$

where c is the mean value of m on $\mathbb{S}^{d-1} \subset \mathbb{R}^d$ and the functions m and Ω are related by

$$m(\boldsymbol{\omega}) = c + \int_{\mathbb{S}^{d-1}} \left(\frac{\pi i}{2} \text{sign}(\boldsymbol{\omega} \cdot \mathbf{y}) + \log \left(\frac{1}{|\boldsymbol{\omega} \cdot \mathbf{y}|} \right) \right) \Omega(\mathbf{y}) d\sigma(\mathbf{y}).$$

3 Spherical harmonics and singular integral operators

The class of smooth admissible functions with which we shall be primarily concerned are the spherical harmonics. These can be viewed as a multi-dimensional extension of the trigonometric polynomials in one dimension. The key property is that it is possible to construct an orthonormal basis of $L_2(\mathbb{S}^{d-1})$ using spherical harmonics. As Fourier multipliers, the spherical

harmonics of degree one correspond to the Riesz transform, which has previously been used to construct steerable wavelets [28]. It turns out that transforms based on any of the spherical harmonics have similar properties. One remarkable aspect is the symmetric role assumed by the spatial and frequency variables and the fact that the Fourier transforms are part of the same family, cf. [26, Chapter IV],[24, Chapter III]. For the benefit of the reader, we review the properties of the spherical harmonics that are relevant for our purpose.

To begin, we recall that a homogeneous polynomial of degree ℓ on \mathbb{R}^d is a linear combination of monomials

$$\sum_{|\alpha|=\ell} c_{\alpha} \omega^{\alpha},$$

where the degree $|\alpha| = \alpha_1 + \dots + \alpha_d$ of each monomial $\omega^{\alpha} = \omega_1^{\alpha_1} \dots \omega_d^{\alpha_d}$ is ℓ and each c_{α} is complex. A homogeneous polynomial P is harmonic if it satisfies Laplace's equation

$$0 = \Delta P = \sum_{k=1}^d \left(\frac{dP}{d\omega_k} \right)^2.$$

The spherical harmonics of degree ℓ , denoted by \mathcal{H}_{ℓ} , can then be defined as the restrictions to the sphere of such polynomials. Specifically, every homogeneous harmonic polynomial P of degree ℓ defines a spherical harmonic $Y \in \mathcal{H}_{\ell}$ by the equation

$$Y(\omega) = P(\omega)$$

for any $\omega \in \mathbb{S}^{d-1} \subset \mathbb{R}^d$. Moreover, any homogeneous polynomial P of degree $k \geq 0$ can be expanded as

$$P(\omega) = P_0(\omega) + \dots + |\omega|^{2\ell} P_{\ell}(\omega),$$

where each P_j is a homogeneous harmonic polynomial of degree $k - 2j$, cf. [26, Theorem IV.2.1]. This means that the proposed steerable wavelet construction includes previous constructions that utilized higher-order Riesz transforms.

An important property that is implicit in many results is that spherical harmonics of different degrees are orthogonal on the sphere \mathbb{S}^{d-1} , cf. [14, Lemma 2] or [26, Corollary IV.2.4]. In particular, since P_0 is constant, this

implies the spherical harmonics of positive degree have mean zero on the sphere:

$$\int_{\mathbb{S}^{d-1}} Y(\boldsymbol{\omega}) d\sigma(\boldsymbol{\omega}) = 0$$

for any $Y \in \mathcal{H}_\ell$ with $\ell > 0$.

The dimension of \mathcal{H}_ℓ can be computed to be

$$N(d, \ell) = \binom{d + \ell - 1}{\ell} - \binom{d + \ell - 3}{\ell - 2}.$$

We shall use the following notation to denote a real-valued orthonormal basis of \mathcal{H}_ℓ :

$$\left\{ Y_{\ell,k} : \mathbb{S}^{d-1} \rightarrow \mathbb{C} : k = 1, \dots, N(d, \ell) \right\}.$$

Summing over $\ell \leq \ell_{\max}$, we can determine the dimension of the space of spherical harmonics of degree at most ℓ_{\max} to be $N(d + 1, \ell_{\max})$, cf. [14, p. 4] or [31, Chapter 17]. Explicit constructions of these basis functions are known. For example, the three-dimensional case is analyzed in detail in [2, Chapter 9], and a general approach for higher dimensions is given in [15, Chapter 2].

We conclude this review with the property that is most important for our purpose: any orthonormal basis $\{Y_{\ell,k}\}$ of \mathcal{H}_ℓ satisfies

$$\sum_k \frac{\sigma(\mathbb{S}^{d-1})}{N(d, \ell)} |Y_{\ell,k}(\boldsymbol{\omega})|^2 = 1, \quad (3)$$

cf. [14, Theorem 2] or [26, Corollary IV.2.9]. This is indeed a powerful property, for it implies that the collection of Fourier multipliers

$$\left\{ m_k(\boldsymbol{\omega}) = \sqrt{\frac{\sigma(\mathbb{S}^{d-1})}{N(d, \ell)}} Y_{\ell,k} \left(\frac{\boldsymbol{\omega}}{|\boldsymbol{\omega}|} \right) : k = 1, \dots, N(d, \ell) \right\}$$

is admissible. Furthermore, as we shall see shortly, this property allows us to define a collection of admissible multipliers which contains all of the spherical harmonics up to a fixed degree ℓ_{\max} .

3.1 Harmonic Riesz transforms

The motivation for steerable wavelets is to provide a wavelet decomposition which more accurately accounts for local orientation of data. In a series

of papers [27, 28, 29], steerable wavelets have been constructed using the Riesz transform and its higher-order variants, taking advantage of their scale and rotation invariance and their unitary character. We recall that the Fourier multiplier of the order ℓ Riesz transform of an $L_2(\mathbb{R}^d)$ function f is a vector valued function whose components are $(\sqrt{\ell!/\alpha!})\omega^\alpha/|\omega|^\ell$, where $|\alpha| = \ell$. In the spatial domain, this translates into a principal value singular integral. The use of spherical harmonics generalizes but also simplifies this construction, thanks to its orthogonality properties.

To make our construction precise, we designate the Fourier multiplier transforms associated with the spherical harmonics as harmonic Riesz transforms.

Definition 3.1. For any positive integer ℓ_{\max} and any unit vector $\mathbf{c} = (c_0, \dots, c_{\ell_{\max}}) \in \mathbb{R}^{\ell_{\max}+1}$, we define the order ℓ_{\max} harmonic Riesz transform to be the multiplier transform $T_{\mathcal{M}}$, where

$$\mathcal{M} = \left\{ m_{\ell,k}(\omega) = c_\ell \sqrt{\frac{\sigma(\mathbb{S}^{d-1})}{N(d,\ell)}} Y_{\ell,k} \left(\frac{\omega}{|\omega|} \right) \right\},$$

where ℓ ranges from 0 to ℓ_{\max} and k ranges from 1 to $N(d,\ell)$.

Note that if \mathbf{c} contains entries which are zero, the corresponding multipliers are not included in the transform. The admissibility of this transform follows immediately from Equation (3). Furthermore, when applied to smooth functions with vanishing moments, the transform preserves decay as well as vanishing moments.

Theorem 3.2. *Let ψ be a differentiable function (or wavelet) with vanishing moments of order $N \geq 1$ such that ψ and its derivatives satisfy the decay estimates*

1. $|\psi(\mathbf{x})| \leq C(1 + |\mathbf{x}|)^{-d-N+\epsilon}$,
2. $|D^\alpha \psi(\mathbf{x})| \leq C(1 + |\mathbf{x}|)^{-d-N-1+\epsilon}$, $|\alpha| = 1$

for some $C > 0$ and $0 \leq \epsilon < 1$. Then for any $\ell \geq 0$ and any $1 \leq k \leq N(d,\ell)$, the corresponding component of any harmonic Riesz transform $T_{\mathcal{M}}(\psi)$ has decay similar to ψ and maintains the same number of vanishing moments, i.e.

$$|[T_{\mathcal{M}}(\psi)]_{\ell,k}(\mathbf{x})| \leq C(1 + |\mathbf{x}|)^{-d-N+\epsilon'}$$

for some $0 \leq \epsilon' < 1$ and $[T_{\mathcal{M}}(\psi)]_{\ell,k}$ has N vanishing moments.

Proof. This follows from the proofs of [30, Theorems 3.2 and 3.4]. Those results were stated for the first-order Riesz transform; however, one can verify that they hold for a more general class of principal value singular integral operators.

The kernels of the singular integral operators that define the Riesz transform are $K(\mathbf{x}) = x_k/|\mathbf{x}|^{d+1}$, while the kernels used in the harmonic Riesz transforms have the form $\tilde{K} = P(\mathbf{x})/|\mathbf{x}|^{d+\ell}$, where P is a homogeneous harmonic polynomial of degree ℓ . The essential properties of Riesz kernels that were used in the proof of those theorems were:

- K has mean zero on the unit sphere;
- the kernel K and its derivatives satisfy certain decay conditions;
- K is smooth away from the origin, so that it can be well approximated locally by polynomials.

It can be verified that each of these conditions holds for \tilde{K} , and hence the results of [30] are applicable as well. \square

3.2 Generalized harmonic Riesz transforms

While directly using the spherical harmonics in a steerable wavelet frame provides a means of categorizing data, we would like to extend this method to make our wavelets adaptable and possibly easier to implement. The approach we take is to compose the harmonic Riesz transforms with matrices representing isometries. Using an isometry, the derived collection of Fourier multipliers will again be admissible.

Definition 3.3. Let $T_{\mathcal{M}}$ be a harmonic Riesz transform consisting of N elements. Additionally, let \mathbf{U} be a complex valued matrix of size $n_{\max} \times N$, which represents an isometry; i.e., $\mathbf{U}^T \mathbf{U}$ is the identity matrix of size N . We define the associated generalized harmonic Riesz transform of $f \in L_2(\mathbb{R}^d)$ to be the vector-valued function $T_{\mathcal{M}, \mathbf{U}}(f)$ obtained by applying \mathbf{U} to the harmonic Riesz transform $T_{\mathcal{M}}(f)$ of f ; i.e., the components of $T_{\mathcal{M}, \mathbf{U}}(f)$ are linear combinations of the components of $T_{\mathcal{M}}(f)$.

Note that in this definition, we deal with two admissible families. Therefore, we have denoted the size of the original family as N , in order to reserve n_{\max} for the size of the derived family. Also, note that a consequence of the isometry condition is that $n_{\max} \geq N$. Additionally, if the initial harmonic Riesz transform $T_{\mathcal{M}}$ is defined by a vector \mathbf{c} with all non-zero entries, then in the above definition $N = N(d+1, \ell_{\max})$.

4 Steerable wavelets

In previous sections we covered the mathematical tools necessary to transform a wavelet frame into a steerable one. In this section, we complete the construction by introducing an appropriate primal wavelet basis. Our choice is the direct extension of the two dimensional case [27].

Proposition 4.1. *Let $h : [0, \infty) \rightarrow \mathbb{R}$ be a smooth function satisfying:*

- (1) $h(\omega) = 0$ for $|\omega| > 1/2$
- (2) $\sum_{j \in \mathbb{Z}} |h(2^j \omega)|^2 = 1$
- (3) $\left. \frac{d^n h(\omega)}{d\omega^n} \right|_{\omega=0} = 0$ for $n = 0, \dots, N$.

Then the isotropic mother wavelet ψ whose d -dimensional Fourier transform is given by

$$\widehat{\psi}(\boldsymbol{\omega}) = h(|\boldsymbol{\omega}|)$$

generates a tight wavelet frame of $L_2(\mathbb{R}^d)$ whose basis functions

$$\psi_{j,\mathbf{k}}(\mathbf{x}) = \psi_j(\mathbf{x} - 2^j \mathbf{k}) \quad \text{with} \quad \psi_j(\mathbf{x}) = 2^{-jd/2} \psi(2^{-j} \mathbf{x})$$

are isotropic with vanishing moments up to order N . Additionally, any $L_2(\mathbb{R}^d)$ function f can be represented as

$$f = \sum_{j \in \mathbb{Z}} \sum_{\mathbf{k} \in \mathbb{Z}^d} \langle f, \psi_{j,\mathbf{k}} \rangle \psi_{j,\mathbf{k}}.$$

Proof. This follows from a combination of Parseval's identity for Fourier transforms and Plancherel's identity for Fourier series. \square

We now define a steerable wavelet frame to be a generalized harmonic Riesz transform of a primal isotropic frame. Theorem 2.4 guarantees that the frame bounds are preserved and that we maintain the reproduction property:

$$\forall f \in L_2(\mathbb{R}^d), \quad f = \sum_{j \in \mathbb{Z}} \sum_{\mathbf{k} \in \mathbb{Z}^d} \sum_{n=1}^{n_{\max}} \langle f, [T_{\mathcal{M}, \mathbf{U}} \psi_{j,\mathbf{k}}]_n \rangle [T_{\mathcal{M}, \mathbf{U}} \psi_{j,\mathbf{k}}]_n$$

As we are applying the harmonic Riesz transforms to isotropic functions, the transform can be reduced to a more manageable form. Essentially, the following is a simplification of Theorem 2.5, which uses radial symmetry to reduce the Fourier transform to a one dimensional integral.

Theorem 4.2. [26, Theorem IV.3.10] Suppose $d \geq 2$ and $\widehat{\psi} \in L^2(\mathbb{R}^d) \cap L^1(\mathbb{R}^d)$ has the form

$$\widehat{\psi}(\boldsymbol{\omega}) = h(|\boldsymbol{\omega}|) \frac{P(\boldsymbol{\omega})}{|\boldsymbol{\omega}|^\ell},$$

where P is a homogeneous harmonic polynomial of degree ℓ , then ψ has the form $\psi(\mathbf{x}) = F(|\mathbf{x}|)P(\mathbf{x})$ where

$$F(r) = 2\pi i^\ell r^{-(d+2\ell-2)/2} \int_0^\infty h(s) J_{(d+2\ell-2)/2}(2\pi r s) s^{d/2} ds$$

and J_ν is the Bessel function of the first kind of order ν .

5 Directional wavelets using zonal harmonics

We use the term steerable to convey the fact that the wavelets we construct are intended to be rotated (or steered) to provide a better analysis of the data. To see how this is accomplished, let us consider a generic steerable wavelet ψ_{Gen} ; i.e., an element of the generalized harmonic Riesz transform $T_{\mathcal{M}, \mathbf{U}}(\psi)$ where ψ is a primal isotropic wavelet. Each such function is of the form

$$\widehat{\psi}_{\text{Gen}}(\boldsymbol{\omega}) = h(|\boldsymbol{\omega}|) \sum_{\ell=0}^{\ell_{\max}} \sum_{k=1}^{N(d,\ell)} u_{\ell,k} Y_{\ell,k} \left(\frac{\boldsymbol{\omega}}{|\boldsymbol{\omega}|} \right)$$

where the coefficients $u_{\ell,k}$ are related to the rows of \mathbf{U} . Of particular interest are the cases where

$$u_{\ell,k} = u_\ell Y_{\ell,k} \left(\frac{\boldsymbol{\omega}_0}{|\boldsymbol{\omega}_0|} \right)$$

for some $\boldsymbol{\omega}_0 \in \mathbb{R}^d \setminus \{0\}$. The resulting spherical function is a zonal function, and it has the form

$$\sum_{\ell=0}^{\ell_{\max}} u_\ell \frac{N(d,\ell)}{\sigma(\mathbb{S}^{d-1})} P_\ell \left(d; \frac{\boldsymbol{\omega}_0}{|\boldsymbol{\omega}_0|} \cdot \frac{\boldsymbol{\omega}}{|\boldsymbol{\omega}|} \right)$$

where $P_\ell(d; \cdot)$ is a generalized Legendre polynomial of degree ℓ , cf. [15, Section 1.2]. A formula for these polynomials is

$$P_\ell(d; x) = \ell! \Gamma \left(\frac{d-1}{2} \right) \sum_{l=0}^{\lfloor \ell/2 \rfloor} \left(\frac{-1}{4} \right)^l \frac{(1-x^2)^l x^{\ell-2l}}{l!(\ell-2l)! \Gamma(l + (d-1)/2)}.$$

for $x \in [-1, 1]$.

One benefit of using these zonal functions is that we can compute rotations fairly effortlessly using the formula:

$$P_\ell(d; \mathbf{R}\boldsymbol{\omega}_0 \cdot \mathbf{R}\boldsymbol{\omega}) = P_\ell(d; \boldsymbol{\omega}_0 \cdot \boldsymbol{\omega}),$$

for any matrix \mathbf{R} satisfying $\mathbf{R}^{-1} = \mathbf{R}^T$. In addition to making rotations straightforward, this structure implies that the value of $\widehat{\psi}_{\text{Gen}}$ is determined solely by the distance of $\boldsymbol{\omega}_0/|\boldsymbol{\omega}_0|$ from $\boldsymbol{\omega}/|\boldsymbol{\omega}|$ on the sphere; i.e., it is a zonal function.

Zonal functions have proved to be particularly valuable for approximation on spheres. For example, positive linear combinations of generalized Legendre polynomials are positive semi-definite functions, which can be used for interpolation [19, 20, 31, 32]. Indeed, the steerable wavelet construction we propose uses polynomials which are positive semi-definite on the sphere.

5.1 Two dimensions

We would now like to present some information regarding the implementation of steerable wavelets based on generalized Legendre polynomials, and we shall start with a two-dimensional motivating example. A basis for the circular harmonics of degree ℓ on \mathbb{S}^1 is

$$\{\sin(\ell\omega), \cos(\ell\omega)\}.$$

Considering Definition 3.3, we shall define an isometry to generate a new partition of unity. Note that we shall neglect certain scaling factors, as they do not impact the underlying principle.

Given a point $\omega_0 \in \mathbb{S}^1$, we define the kernel P in the span of the degree ℓ basis by

$$\begin{aligned} P(\omega_0, \omega) &= \sin(\ell\omega_0) \sin(\ell\omega) + \cos(\ell\omega_0) \cos(\ell\omega) \\ &= \cos(\ell(\omega_0 - \omega)). \end{aligned}$$

For a collection of points $\{\omega_n\}_{n=1}^{n_{\max}}$ the matrix that transforms the degree ℓ basis $\{\sin(\ell\omega), \cos(\ell\omega)\}$ into $\{\cos(\ell(\omega - \omega_1)), \dots, \cos(\ell(\omega - \omega_{n_{\max}}))\}$ is

$$\mathbf{U}_\ell = \begin{pmatrix} \sin(\ell\omega_1) & \cos(\ell\omega_1) \\ \sin(\ell\omega_2) & \cos(\ell\omega_2) \\ \vdots & \vdots \\ \sin(\ell\omega_{n_{\max}}) & \cos(\ell\omega_{n_{\max}}) \end{pmatrix}.$$

To ensure that the columns of \mathbf{U}_ℓ are orthogonal, we require

$$\begin{aligned} 0 &= \sum_{n=1}^{n_{\max}} \sin(\ell\omega_n) \cos(\ell\omega_n) \\ &= \sum_{n=1}^{n_{\max}} \sin(2\ell\omega_n). \end{aligned}$$

We could consider choosing the points ω_n to be roots of $\sin(2\ell\cdot)$; however, this approach would be less useful in higher dimensions. Instead, we shall use a circular quadrature rule with equal weights. Let $\{\omega_n\}_{n=1}^{n_{\max}}$ be a set of points for which

$$\int_0^{2\pi} p(\omega) d\omega = \frac{2\pi}{n_{\max}} \sum_{n=1}^{n_{\max}} p(\omega_n) \quad (4)$$

for all trigonometric polynomials of degree at most 2ℓ . Then

$$0 = \int_0^{2\pi} \sin(2\ell\omega) d\omega.$$

would imply that the columns of \mathbf{U}_ℓ are orthogonal, and hence that \mathbf{U}_ℓ is an isometry (after normalization).

Sets of points satisfying Equation (4) are well known and are referred to as spherical t -designs, where t indicates the maximum degree polynomial for which quadrature holds. Such sets are known to exist for arbitrarily large t [21], and the most natural choices consist of equidistributed points [1]. Incidentally, Simoncelli's two dimensional equiangular steerable wavelet construction can be reinterpreted in terms of such t -designs [18].

5.2 Higher dimensions

Fix $\ell > 0$ and let $\{Y_{\ell,k}\}_{k=1}^{N(d,\ell)}$ be an orthonormal basis for the spherical harmonics of degree ℓ on \mathbb{S}^{d-1} . As in the two-dimensional case, we shall define an isometry to generate a new partition of unity. Specifically, we select \mathbf{U}_ℓ to be the matrix satisfying

$$[\mathbf{U}_\ell]_{n,k} = Y_{\ell,k}(\omega_n)$$

for a collection of points $X = \{\omega_n\}_{n=1}^{n_{\max}} \subset \mathbb{S}^{d-1}$ to be specified. Delsarte et. al refer to this as the ℓ th characteristic matrix associated with X [4,

Definition 3.4]. Applying \mathbf{U}_ℓ to the basis generates a new basis of zonal polynomials

$$\sum_{k=1}^{N(d,\ell)} Y_{\ell,k}(\boldsymbol{\omega}_m) Y_{\ell,k}(\boldsymbol{\omega}) = \frac{N(d,\ell)}{\sigma(\mathbb{S}^{d-1})} P_\ell(d; \boldsymbol{\omega}_m \cdot \boldsymbol{\omega}).$$

In order to make the columns of \mathbf{U}_ℓ orthogonal, we need

$$0 = \sum_{n=1}^{n_{\max}} Y_{\ell,k}(\boldsymbol{\omega}_n) Y_{\ell,k'}(\boldsymbol{\omega}_n) \quad (5)$$

for $k \neq k'$. As in the two-dimensional case, we choose the points X to form a spherical 2ℓ -design, so that

$$\frac{\sigma(\mathbb{S}^{d-1})}{n_{\max}} \sum_{n=1}^{n_{\max}} p(\boldsymbol{\omega}_n) = \int_{\mathbb{S}^{d-1}} p(\boldsymbol{\omega}) d\sigma(\boldsymbol{\omega}) \quad (6)$$

for spherical harmonics of degree at most 2ℓ . In fact, the conditions on X considered in (5) and (6) are almost equivalent, cf. [4, Remark 5.4]. On \mathbb{S}^2 , examples of t -designs are provided by the vertices of platonic solids: the vertices of an icosahedron or a dodecahedron constitute 4-designs [8]. More generally, t -designs appear as orbits of elements of \mathbb{S}^{d-1} under the action of a finite subgroup of the orthogonal group on the sphere [1]. Some specific examples are given in an online library of t -designs [9]. In particular, this library contains t -designs on \mathbb{S}^2 , where t ranges from 0 to 21.

5.3 Localized kernels

Concerning the construction of zonal basis functions, one final point to address is localization. The reason for this is that well localized functions can be used to more accurately represent the orientation of data. While we would ideally like to use locally supported functions, they cannot be represented as polynomials. Therefore, we shall instead use a normalized polynomial approximation of the identity.

Let us first recall that on the circle, the delta function can be represented by

$$\delta_0(\omega) = \frac{1}{2\pi} + \frac{1}{\pi} \sum_{\ell=1}^{\infty} \cos(\ell\omega),$$

and we could construction an approximation by truncating this series, producing a Dirichlet kernel. The problem with such a construction is that

the Dirichlet kernel is highly oscillatory. Therefore, we instead propose to construct approximate identities analogous to the scaling functions of Freedman et al. [6, Section 11.1.3]. For this construction, we begin with a compactly supported function $\hat{a} : [0, \infty) \rightarrow [0, 1]$ satisfying:

1. $\hat{a}(0) = 1$, $\hat{a}(1) = 0$, and $\hat{a}(\omega) > 0$ for $\omega \in (0, 1)$;
2. \hat{a} is monotonically decreasing;
3. \hat{a} is continuous at 0 and piecewise continuous on $[0, 1]$.

Then an approximation to δ_0 is given by

$$S_{\ell_{\max}}(\omega) = \frac{1}{2\pi} + \frac{1}{\pi} \sum_{\ell=1}^{\ell_{\max}} \hat{a}\left(\frac{\ell}{\ell_{\max} + 1}\right) \cos(\ell\omega).$$

Now, based on the results of the previous section, we know that an admissible zonal basis is given by

$$\Lambda(\omega - \omega_n) = c_0 \sqrt{\frac{1}{n_{\max}}} + \sum_{\ell=1}^{\ell_{\max}} c_\ell \sqrt{\frac{2}{n_{\max}}} \cos(\ell(\omega - \omega_n))$$

for any unit vector \mathbf{c} and any circular $2\ell_{\max}$ -design $\{\omega_n\}_{n=1}^{n_{\max}}$. Therefore we should choose \mathbf{c} to be a normalization of the coefficients of $S_{\ell_{\max}}$ to obtain a well localized basis.

Several examples of such functions appear in the literature. For instance, in [6, Section 11.4.3] the authors propose the cubic polynomial

$$\hat{a}(\omega) = (1 - \omega)^2(1 + 2\omega) \tag{7}$$

to produce a function $S_{\ell_{\max}}(\omega)$ with suppressed oscillations. An alternative choice, based on the localization analysis of [16] and a construction from [11], is to choose \hat{a} to be a B-spline centered at 0. B-splines also appear implicitly in Simoncelli's steerable pyramid construction [18]. The angular part of his wavelets are of the form

$$\begin{aligned} \cos(\omega)^{\ell_{\max}} &= e^{i\omega\ell_{\max}} \sum_{\ell=0}^{\ell_{\max}} \binom{\ell_{\max}}{\ell} e^{-i\omega 2\ell} \\ &= \sum_{\ell=0}^{\ell_{\max}} \binom{\ell_{\max}}{\ell} e^{i\omega(\ell_{\max} - 2\ell)}, \end{aligned}$$

and the binomial coefficients are the discrete analog of the B-splines.

In higher dimensions we can apply the same analysis to construct localized kernels. This leads us to approximate the delta function at $\boldsymbol{\omega}_0$ by

$$S_{\ell_{\max}}(\boldsymbol{\omega}_0 \cdot \boldsymbol{\omega}) = \sum_{\ell=0}^{\ell_{\max}} \hat{a} \left(\frac{\ell}{\ell_{\max} + 1} \right) \frac{N(d, \ell)}{\sigma(\mathbb{S}^{d-1})} P_{\ell}(d; \boldsymbol{\omega}_0 \cdot \boldsymbol{\omega}).$$

Additionally, our analysis requires that we use kernels of the form

$$\Lambda(\boldsymbol{\omega}_0 \cdot \boldsymbol{\omega}) = \sum_{\ell=0}^{\ell_{\max}} c_{\ell} \sqrt{\frac{N(d, \ell)}{n_{\max}}} P_{\ell}(d; \boldsymbol{\omega}_0 \cdot \boldsymbol{\omega});$$

as before, we can adjust the coefficients of $S_{\ell_{\max}}$ to determine an appropriate \mathbf{c} in Λ . Note that the choice of \hat{a} and n_{\max} must be balanced to produce a good kernel. Choosing a smoother \hat{a} produces a kernel with less oscillation; however, its main lobe will be less localized. To compensate, we could increase the degree ℓ_{\max} , but this means that we need to increase n_{\max} ; i.e., utilize a larger collection of basis functions. Example zonal kernels Λ are plotted in Figure 1.

If localization is of primary importance, one can define a variance on the sphere, which should be minimized by a polynomial of a given degree. Indeed, this approach was used in [27] for steerable wavelets on \mathbb{R}^2 . Furthermore, an uncertainty principle for localization in both space and frequency was studied in [6, 7, 17]. Related results are contained in [5], where the author considers polynomials whose degrees lie within a given range. Additional work concerning the localization of spherical Slepian functions can be found in [22, 23].

In the two-dimensional case, Unser et al. extend the definition of variance and allow for a general class of weight functions [27, Appendix A]. A similar generalized definition of variance was introduced by Michel on $\mathbb{S}^2 \subset \mathbb{R}^3$ [12]. However, instead of using Fourier analysis and Bochner's theorem, his results utilize the theory of orthogonal polynomials. In Appendix A, we provide an extension of this type of approach to the higher dimensional setting that can serve as a basis for the design of wavelets with an optimal angular selectivity.

6 Construction and steering

In this section we lay out the construction of the zonal basis and spherical harmonic basis in more detail. Furthermore, we show how the wavelets can be steered using matrix multiplication. To begin the construction, we choose a maximum degree ℓ_{\max} and a real orthonormal basis of spherical harmonics




Figure 1: Plot of zonal kernels $\Lambda(\cos(\cdot))$ for $d = 3$, $\ell_{\max} = 10$, and $n_{\max} = 216$. The kernel on the left was constructed using \hat{a} from (7), while the functions on the right were constructed by choosing \hat{a} to be B-splines. In the plot on the right, the dashed line corresponds to a linear B-spline, and the solid line corresponds to a cubic B-spline.

of degree at most ℓ_{\max} in a vector $[\mathbf{Y}]_m = Y_m$ of length $N(d+1, \ell_{\max})$. Then, given a unit vector $\mathbf{c} = (c_0, \dots, c_{\max}) \in \mathbb{R}^{\ell_{\max}+1}$, we define a diagonal matrix \mathbf{C} of size $N(d+1, \ell_{\max}) \times N(d+1, \ell_{\max})$ as follows: For $[\mathbf{Y}]_m$ of degree ℓ , we set

$$[\mathbf{C}]_{m,m} = c_\ell \sqrt{\frac{\sigma(\mathbb{S}^{d-1})}{N(d, \ell)}}.$$

This implies that the entries of $\mathbf{C}\mathbf{Y}$ form an admissible collection. Next we choose a collection of points on the sphere $X_0 = \{\boldsymbol{\omega}_n\}_{n=1}^{n_{\max}}$, which form a $2\ell_{\max}$ -design on \mathbb{S}^{d-1} . To construct the final admissible collection, we define the $n_{\max} \times N(d+1, \ell_{\max})$ matrix \mathbf{U}_{X_0} as

$$[\mathbf{U}_{X_0}]_{n,m} = \sqrt{\frac{\sigma(\mathbb{S}^{d-1})}{n_{\max}}} Y_m(\boldsymbol{\omega}_n).$$

Our admissible collection is then given by $\mathbf{Z}_{X_0} := \mathbf{U}_{X_0} \mathbf{C}\mathbf{Y}$.

For steering, we use the isometry property of \mathbf{U}_{X_0} . Let \mathbf{R} be a rotation matrix and define X_1 to be the collection of points obtained by rotating the elements of X_0 by \mathbf{R} . We then can expand the elements of the rotated basis \mathbf{Z}_{X_1} in terms of the original basis \mathbf{Z}_{X_0} as

$$\begin{aligned} \mathbf{Z}_{X_1} &= \mathbf{U}_{X_1} \mathbf{C}\mathbf{Y} \\ &= \mathbf{U}_{X_1} \mathbf{U}_{X_0}^T \mathbf{U}_{X_0} \mathbf{C}\mathbf{Y} \\ &= (\mathbf{U}_{X_1} \mathbf{U}_{X_0}^T) \mathbf{Z}_{X_0}. \end{aligned}$$

Since the steerable wavelet frame is a Parseval frame, the steering matrix, which transforms the wavelet coefficients from the original basis into the coefficients of the rotated basis, corresponds to the matrix mapping \mathbf{Z}_{X_0} to \mathbf{Z}_{X_1} . Hence the $n_{\max} \times n_{\max}$ steering matrix is given by $\mathbf{S} := \mathbf{U}_{X_1} \mathbf{U}_{X_0}^T$. Using the properties of zonal harmonics, we can see that the entries of this matrix are pointwise evaluations of a zonal polynomial $\Lambda_{\ell_{\max}}$ (see Figure 2 for an example). Precisely




Figure 2: Plot of the kernel $\Lambda_{\ell_{\max}}(\cos(\cdot))$ for $\ell_{\max} = 10$ and $n_{\max} = 216$.

$$\begin{aligned}
\Lambda_{\ell_{\max}}(\boldsymbol{\omega} \cdot \tilde{\boldsymbol{\omega}}) &= \frac{\sigma(\mathbb{S}^{d-1})}{n_{\max}} \sum_{m=1}^{N(d+1, \ell_{\max})} Y_m(\boldsymbol{\omega}) Y_m(\tilde{\boldsymbol{\omega}}) \\
&= \frac{\sigma(\mathbb{S}^{d-1})}{n_{\max}} \sum_{\ell=0}^{\ell_{\max}} \frac{N(d, \ell)}{\sigma(\mathbb{S}^{d-1})} P_{\ell}(d; \boldsymbol{\omega} \cdot \tilde{\boldsymbol{\omega}}) \\
&= \sum_{\ell=0}^{\ell_{\max}} \frac{N(d, \ell)}{n_{\max}} P_{\ell}(d; \boldsymbol{\omega} \cdot \tilde{\boldsymbol{\omega}})
\end{aligned}$$

and

$$[\mathbf{S}]_{n_1, n_2} = \Lambda_{\ell_{\max}}(\mathbf{R}\boldsymbol{\omega}_{n_1} \cdot \boldsymbol{\omega}_{n_2}).$$

Interestingly, the steering operation is very much akin to an interpolation that uses $\Lambda_{\ell_{\max}}$ as a kernel (cf. Figure 2).

Depending on the number of points, the steering matrix can be quite large. As an alternative, one could work directly with the orthonormal spherical harmonic basis. In this situation, the steering matrix reduces significantly as it is a block diagonal matrix with blocks of size $N(d, \ell)$. To see this, let us define the vector \mathbf{Y} by

$$\mathbf{Y} = \begin{pmatrix} \mathbf{Y}_0 \\ \mathbf{Y}_1 \\ \vdots \\ \mathbf{Y}_{\ell_{\max}} \end{pmatrix},$$

where each \mathbf{Y}_ℓ is a vector whose components are an orthonormal basis of the spherical harmonics of degree ℓ :

$$\mathbf{Y}_\ell = \begin{pmatrix} Y_{\ell,1}(\boldsymbol{\omega}) \\ Y_{\ell,2}(\boldsymbol{\omega}) \\ \vdots \\ Y_{\ell,N(d,\ell)}(\boldsymbol{\omega}) \end{pmatrix}.$$

In order to make the entries of \mathbf{Y} an admissible collection, we multiply by the block diagonal matrix

$$\mathbf{C} = \begin{pmatrix} \mathbf{C}_0 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{C}_1 & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{C}_{\ell_{\max}} \end{pmatrix},$$

where \mathbf{C}_ℓ is the $N(d,\ell) \times N(d,\ell)$ diagonal matrix with entries

$$[\mathbf{C}_\ell]_{k,k} = \sqrt{\frac{\sigma(\mathbb{S}^{d-1})}{(\ell_{\max} + 1)N(d,\ell)}}.$$

With \mathbf{C} defined in this way, we are giving equal weight to each degree ℓ . One could choose an alternative weighting, but it must be constant over a given degree for the partition of unity property to hold. Now to construct a steering matrix, we need to find an expression for the spherical harmonics in terms of any rotation of them. Orthogonality between degrees of spherical harmonics makes this especially nice because it means that we can rotate any given degree independently of the others. Therefore, let us fix ℓ and consider the problem of steering the functions in $\mathbf{C}_\ell \mathbf{Y}_\ell$. As \mathbf{C}_ℓ is a constant multiple of the identity matrix, this is equivalent to finding a steering matrix for \mathbf{Y}_ℓ . Since we are dealing with an orthonormal basis, we can expand any Y_{ℓ,k_0} as

$$Y_{\ell,k_0}(\mathbf{R}\boldsymbol{\omega}) = \sum_{k=1}^{N(d,\ell)} \langle Y_{\ell,k_0}(\mathbf{R}\cdot), Y_{\ell,k} \rangle Y_{\ell,k}(\boldsymbol{\omega}),$$

and computation of these inner products can be handled using a quadrature rule. Specifically, for any 2ℓ -design $\{\omega_n\}_{n=1}^N$, we have

$$\langle Y_{\ell,k_0}(\mathbf{R}\cdot), Y_{\ell,k} \rangle = \frac{\sigma(\mathbb{S}^{d-1})}{N} \sum_{n=1}^N Y_{\ell,k_0}(\mathbf{R}\omega_n) Y_{\ell,k}(\omega_n).$$

Therefore, the $N(d, \ell) \times N(d, \ell)$ matrix \mathbf{V}_ℓ with entries

$$[\mathbf{V}_\ell]_{k,k'} = \frac{\sigma(\mathbb{S}^{d-1})}{N} \sum_{n=1}^N Y_{\ell,k'}(\mathbf{R}\omega_n) Y_{\ell,k}(\omega_n)$$

transforms the original basis \mathbf{Y}_ℓ into the rotated basis $\mathbf{Y}_\ell^{\mathbf{R}}$, i.e. $\mathbf{V}_\ell \mathbf{Y}_\ell = \mathbf{Y}_\ell^{\mathbf{R}}$. Consequently, the steering matrix which transforms the wavelet coefficients corresponding to \mathbf{Y}_ℓ into the coefficients corresponding to $\mathbf{Y}_\ell^{\mathbf{R}}$ is also given by \mathbf{V}_ℓ . Note that in two dimensions $N(d, \ell) = 2$ for $\ell \geq 1$, and in three dimensions $N(d, \ell) = 2\ell + 1$. Therefore, these steering matrices can be significantly smaller than the ones used in the zonal construction when a large ℓ_{\max} is chosen.

7 Conclusion and practical summary

Throughout the course of this paper, we have developed the theory of steerable wavelets in any number of dimensions greater than one. The previous two sections were devoted to several technical aspects of implementation, and here we summarize the construction. Since the zonal construction is perhaps more tractable conceptually, we now concentrate on this. One can think of this construction as a generalization for $d > 2$ of Simoncelli's equiangular design.

The substructure of a steerable wavelet frame is an isotropic mother wavelet ψ , which satisfies the conditions of Proposition 4.1. A new, expanded frame is produced by the collection of mother wavelets

$$\left\{ \mathcal{F}^{-1} \{ m_n \hat{\psi} \} : m_n \in \mathcal{M} \right\},$$

where \mathcal{M} is an admissible class of functions. Now given a maximum degree ℓ_{\max} , a unit vector $\mathbf{c} = (c_0, \dots, c_{\ell_{\max}}) \in \mathbb{R}^{\ell_{\max}+1}$, and a spherical $2\ell_{\max}$ -design $X = \{\omega_n\}_{n=1}^{n_{\max}}$, we define the admissible collection

$$\mathcal{M} = \left\{ m_n(\omega) = \sum_{\ell=0}^{\ell_{\max}} c_\ell \sqrt{\frac{N(d, \ell)}{n_{\max}}} P_\ell \left(d; \frac{\omega_n}{|\omega_n|} \cdot \frac{\omega}{|\omega|} \right) : n = 1, \dots, n_{\max} \right\},$$

where the $P_\ell(d; \cdot)$ are generalized Legendre polynomials. The main properties of this basis are summarized in the following theorem.

Theorem 7.1. *The wavelet construction given above defines a tight curvelet-like wavelet frame of $L_2(\mathbb{R}^d)$, where all wavelets are rotated versions of a single wavelet template (per scale). Furthermore, these wavelets are parametrized with a set of coefficients c_ℓ that can be chosen arbitrarily.*

Ideally, the vector \mathbf{c} is chosen so that the functions m_n are well localized and peaked at $\boldsymbol{\omega} = \boldsymbol{\omega}_n$. For large collections of points $X = \{\boldsymbol{\omega}_n\}_{n=1}^{n_{\max}}$ on the sphere and well-localized functions $m_n(\boldsymbol{\omega})$, this new frame can detect the orientation of data and provide more information concerning structure. However, one can also work with a smaller collection \mathcal{M} and take advantage of the steering property to orient the wavelet frame in a data-adaptive fashion. In order to steer the wavelet basis, we need a matrix which transforms the wavelet coefficients upon rotation of the wavelet basis. The structure of the zonal basis allows us to use the $n_{\max} \times n_{\max}$ steering matrix

$$[\mathbf{S}]_{n_1, n_2} = \Lambda_{\ell_{\max}}(\mathbf{R}\boldsymbol{\omega}_{n_1} \cdot \boldsymbol{\omega}_{n_2})$$

for a rotation \mathbf{R} , where

$$\Lambda_{\ell_{\max}}(\boldsymbol{\omega} \cdot \tilde{\boldsymbol{\omega}}) = \sum_{\ell=0}^{\ell_{\max}} \frac{N(d, \ell)}{n_{\max}} P_\ell(d; \boldsymbol{\omega} \cdot \tilde{\boldsymbol{\omega}}).$$

Notice that the entire construction is in terms of the zonal basis, which means we can avoid working directly with the spherical harmonics. On the other hand, working with the spherical harmonics would allow for smaller steering matrices.

A Optimal spherical polynomial constructions

Using the work of Michel [12] as a starting point, in this appendix we shall show how orthogonal polynomials can be used to construct spherical polynomials that minimize energy functionals on \mathbb{S}^{d-1} . In order to simplify computations let us fix a point $\boldsymbol{\omega}_0 \in \mathbb{S}^{d-1}$ and assume that each of the generalized Legendre polynomials in $\{P_\ell(d; \boldsymbol{\omega}_0 \cdot \boldsymbol{\omega})\}_{\ell=0}^{\ell_{\max}}$ has been normalized to have $L_2(\mathbb{S}^{d-1})$ norm one.

Considering the framework of our problem, we shall address the problem of finding the polynomial (centered at $\boldsymbol{\omega}_0$)

$$P_{\mathbf{c}}(\boldsymbol{\omega}_0 \cdot \boldsymbol{\omega}) = \sum_{\ell=0}^{\ell_{\max}} c_{\ell} P_{\ell}(d; \boldsymbol{\omega}_0 \cdot \boldsymbol{\omega})$$

of norm one in $L_2(\mathbb{S}^{d-1})$ that minimizes (or maximizes) an energy functional of the form

$$E(P_{\mathbf{c}}; W) = \int_{\mathbb{S}^{d-1}} |P_{\mathbf{c}}(\boldsymbol{\omega}_0 \cdot \boldsymbol{\omega})|^2 W(\boldsymbol{\omega}_0 \cdot \boldsymbol{\omega}) d\sigma(\boldsymbol{\omega}),$$

where W is an arbitrary positive continuous function. One possibility would be to choose $W(t) = \arccos(t)^2$, so that the energy functional would be

$$\begin{aligned} E(P_{\mathbf{c}}; \arccos(\cdot)^2) &= \int_{\mathbb{S}^{d-1}} |P_{\mathbf{c}}(\boldsymbol{\omega}_0 \cdot \boldsymbol{\omega})|^2 \arccos(\boldsymbol{\omega}_0 \cdot \boldsymbol{\omega})^2 d\sigma(\boldsymbol{\omega}) \\ &= \int_{\mathbb{S}^{d-1}} |P_{\mathbf{c}}(\boldsymbol{\omega}_0 \cdot \boldsymbol{\omega})|^2 \text{dist}(\boldsymbol{\omega}_0, \boldsymbol{\omega})^2 d\sigma(\boldsymbol{\omega}), \end{aligned}$$

where *dist* refers to the spherical distance.

As the energy functional contains only zonal functions, it can be reduced to a simpler form:

$$\begin{aligned} E(P_{\mathbf{c}}; W) &= \frac{\Gamma(\frac{d}{2})}{\sqrt{\pi}\Gamma(\frac{d-1}{2})} \int_0^{\pi} |P_{\mathbf{c}}(\cos(\theta))|^2 W(\cos(\theta)) \sin^{d-2}(\theta) d\theta \\ &= \frac{\Gamma(\frac{d}{2})}{\sqrt{\pi}\Gamma(\frac{d-1}{2})} \int_{-1}^1 |P_{\mathbf{c}}(t)|^2 W(t) (1-t^2)^{(d-3)/2} dt. \end{aligned}$$

Now notice that any polynomial of degree at most ℓ_{\max} is a linear combination of generalized Legendre polynomials $\{P_{\ell}(d; \cdot) : \ell = 0 \dots, \ell_{\max}\}$. Also, the assumptions on W imply that there exists a sequence of polynomials $\{Q_{\ell}\}_{\ell=0}^{\ell_{\max}}$ that are orthonormal with respect to E ,

$$\frac{\Gamma(\frac{d}{2})}{\sqrt{\pi}\Gamma(\frac{d-1}{2})} \int_{-1}^1 Q_{\ell}(t) Q_{\ell'}(t) W(t) (1-t^2)^{(d-3)/2} dt = \delta_{\ell, \ell'},$$

where each Q_{ℓ} is a polynomial of degree ℓ , cf. [10]. Hence, there exists an invertible change of basis matrix from $\{P_{\ell}(d; \cdot)\}_{\ell=0}^{\ell_{\max}}$ to $\{Q_{\ell}\}_{\ell=0}^{\ell_{\max}}$. Consequently, the polynomial $P_{\mathbf{c}}$ that minimizes the functional $E(P_{\mathbf{c}}; W)$ is determined by setting \mathbf{c} to be the unit eigenvector corresponding to the minimal eigenvalue (in absolute value) of the change of basis matrix.

References

- [1] Ei. Bannai and Et. Bannai. A survey on spherical designs and algebraic combinatorics on spheres. *Eur. J. Combin.*, 30(6):1392–1425, 2009.
- [2] R. Beals and R. Wong. *Special functions*, volume 126 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2010.
- [3] A. P. Calderón and A. Zygmund. Singular integral operators and differential equations. *Am. J. Math.*, 79:901–921, 1957.
- [4] P. Delsarte, J. M. Goethals, and J. J. Seidel. Spherical codes and designs. *Geometriae Dedicata*, 6(3):363–388, 1977.
- [5] N. L. Fernández. Optimally space-localized band-limited wavelets on \mathbb{S}^{q-1} . *J. Comput. Appl. Math.*, 199(1):68–79, 2007.
- [6] W. Freeden, T. Gervens, and M. Schreiner. *Constructive approximation on the sphere*. Numerical Mathematics and Scientific Computation. The Clarendon Press Oxford University Press, New York, 1998.
- [7] W. Freeden and V. Michel. Constructive approximation and numerical methods in geodetic research today – an attempt at a categorization based on an uncertainty principle. *J. Geodesy*, 73:452–465, 1999.
- [8] R. H. Hardin and N. J. A. Sloane. New spherical 4-designs. *Discrete Math.*, 106/107:255–264, 1992.
- [9] R. H. Hardin and N. J. A. Sloane. Spherical designs, <http://neilsloane.com/sphdesigns/index.html>. Last Accessed 28 Dec 2012.
- [10] D. Jackson. Series of orthogonal polynomials. *Ann. of Math. (2)*, 34(3):527–545, 1933.
- [11] H. N. Mhaskar, F. J. Narcowich, J. Prestin, and J. D. Ward. Polynomial frames on the sphere. *Adv. Comput. Math.*, 13(4):387–403, 2000.
- [12] V. Michel. Optimally localized approximate identities on the 2-sphere. *Numer. Func. Anal. Opt.*, 32(8):877–903, 2011.
- [13] S. G. Mikhlin. *Multidimensional singular integrals and integral equations*. Translated from the Russian by W. J. A. Whyte. Translation edited by I. N. Sneddon. Pergamon Press, Oxford, 1965.

- [14] C. Müller. *Spherical harmonics*, volume 17 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1966.
- [15] C. Müller. *Analysis of spherical symmetries in Euclidean spaces*, volume 129 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 1998.
- [16] F. J. Narcowich, P. Petrushev, and J. D. Ward. Localized tight frames on spheres. *SIAM J. Math. Anal.*, 38(2):574–594 (electronic), 2006.
- [17] F. J. Narcowich and J. D. Ward. Nonstationary wavelets on the m -sphere for scattered data. *Appl. Comput. Harmon. A.*, 3(4):324–336, 1996.
- [18] J. Portilla and E. P. Simoncelli. A parametric texture model based on joint statistics of complex wavelet coefficients. *Int. J. Comput. Vision*, 40:49–70, 2000.
- [19] I. J. Schoenberg. Positive definite functions on spheres. *Duke Math. J.*, 9:96–108, 1942.
- [20] M. Schreiner. On a new condition for strictly positive definite functions on spheres. *P. Am. Math. Soc.*, 125(2):531–540, 1997.
- [21] P. D. Seymour and T. Zaslavsky. Averaging sets: a generalization of mean values and spherical designs. *Adv. Math.*, 52(3):213–240, 1984.
- [22] F. J. Simons. Slepian functions and their use in signal estimation and spectral analysis. In Willi Freeden, M.Zuhair Nashed, and Thomas Sonar, editors, *Handbook of Geomathematics*, pages 891–923. Springer Berlin Heidelberg, 2010.
- [23] F. J. Simons, F. A. Dahlen, and M. A. Wieczorek. Spatiospectral concentration on a sphere. *SIAM Rev.*, 48(3):504–536 (electronic), 2006.
- [24] E. M. Stein. *Singular integrals and differentiability properties of functions*, volume 30 of *Princeton Mathematical Series*. Princeton University Press, Princeton, N.J., 1970.
- [25] E. M. Stein. *Harmonic analysis: real-variable methods, orthogonality, and oscillatory integrals*, volume 43 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 1993.
- [26] E. M. Stein and G. Weiss. *Introduction to Fourier analysis on Euclidean spaces*, volume 32 of *Princeton Mathematical Series*. Princeton University Press, Princeton, N.J., 1971.

- [27] M. Unser and N. Chenouard. A unifying parametric framework for 2D steerable wavelet transforms. *SIAM J. Imaging Sci.*, in press.
- [28] M. Unser, N. Chenouard, and D. Van De Ville. Steerable pyramids and tight wavelet frames in $L_2(\mathbf{R}^d)$. *IEEE T. Image Process.*, 20(10):2705–2721, 2011.
- [29] M. Unser and D. Van De Ville. Wavelet steerability and the higher-order riesz transform. *IEEE T. Image Process.*, 19(3):636–652, 2010.
- [30] J. P. Ward, K. N. Chaudhury, and M. Unser. Decay properties of Riesz transforms and steerable wavelets. arXiv:1301.2525.
- [31] H. Wendland. *Scattered data approximation*, volume 17 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, Cambridge, 2005.
- [32] Y. Xu and E. W. Cheney. Strictly positive definite functions on spheres. *P. Am. Math. Soc.*, 116(4):977–981, 1992.